Local Copies of OpenAI GPT-4 Are Not Being Distributed for Local Civil Defense Due to OpenAI's Shift Toward Profit

Douglas C. Youvan

doug@youvan.com

September 21, 2024

In the rapidly evolving landscape of artificial intelligence, AI models like OpenAI's GPT-4 hold immense potential for enhancing civil defense, especially during crises such as natural disasters, pandemics, or armed conflicts. However, the decision by OpenAI to shift from its original open-source mission toward a profit-driven model has created significant barriers to the widespread deployment of these powerful tools. By restricting access to GPT-4 through cloud-based services and high subscription fees, OpenAI has effectively prevented local communities and organizations from utilizing this technology for critical civil defense applications. This paper explores the ethical and strategic implications of OpenAI's shift toward profit, highlights the missed opportunities for local AI deployment in crisis scenarios, and proposes solutions to ensure that AI remains accessible for the public good, particularly in high-stakes situations where lives are at risk.

Keywords: OpenAI, GPT-4, artificial intelligence, civil defense, open-source AI, profit-driven model, local deployment, disaster preparedness, ethical AI, cloud-based services, community resilience, crisis management, AI access, decentralized AI, public good. 42 pages.

Introduction

OpenAI was founded in 2015 with the ambitious goal of advancing artificial intelligence in a way that would benefit all of humanity. Its mission was clear: to ensure that AI technology, especially highly advanced forms like artificial general intelligence (AGI), would be developed safely, ethically, and in an open manner so that the benefits could be broadly shared across society. At its core, OpenAI sought to create a future where AI could address global challenges such as poverty, climate change, and healthcare, while preventing the technology from being monopolized by a few powerful entities or misused for harmful purposes.

One of the foundational promises of OpenAI was its commitment to open-source principles. The organization initially positioned itself as a leader in democratizing access to cutting-edge AI technologies. OpenAI's early statements indicated a commitment to sharing research, software, and insights freely with the public. This approach was meant to catalyze innovation in AI by allowing developers, researchers, and organizations around the world to collaborate, build on each other's work, and create AI solutions that served local and global needs. Importantly, this also meant that AI would not remain the exclusive domain of large corporations or governments; instead, small communities, nonprofits, and individual innovators could tap into this powerful technology to solve local challenges.

The promise of open-source AI held particular significance for civil defense applications. In times of crisis—whether due to natural disasters, conflict, or other emergencies—access to powerful AI models could provide vital support. Local communities, especially those in remote or under-resourced areas, could benefit from AI-powered decision-making tools, real-time disaster response guidance, medical support, and resource allocation. Open-source AI would empower them to deploy these models on local hardware without relying on centralized systems, which may be unavailable in a crisis. The vision was clear: by making AI accessible to all, OpenAI could become a critical tool for resilience and safety worldwide.

However, as OpenAI evolved, it shifted from its initial open-source ethos toward a more profit-driven model. In 2019, OpenAI restructured into a capped-profit entity, allowing for outside investments while capping the return investors could make. This marked a fundamental shift in its approach. With the introduction of

paid access models and licensing restrictions, OpenAI's most advanced tools, like GPT-4, were no longer freely accessible to the public. Instead of enabling widespread local deployment, OpenAI's financial model now centers around cloud-based access, where users must pay for API usage. This has created significant barriers for communities that could otherwise benefit from localized AI solutions.

The transition from open-source to a profit-driven model has had direct consequences for civil defense. While the hardware required to run AI models like GPT-4 locally is now affordable for many communities and individuals, the restrictive financial policies around accessing the model itself have kept this technology out of reach for most. The promise of democratized access to AI has been compromised, limiting the potential for communities to deploy these tools in times of crisis.

In this paper, we argue that OpenAI's shift toward profit-driven models has restricted the distribution of local AI tools for civil defense, despite the growing need for decentralized, resilient systems in an increasingly uncertain world. By prioritizing profit over accessibility, OpenAI has undermined its original mission, leaving many communities vulnerable to crises that could otherwise be mitigated through local AI deployment.

The Original Vision: Open-Source for the Public Good

From its inception, OpenAI was founded on the principle of developing artificial intelligence as an open-source resource, ensuring that its benefits could be shared across all of humanity, not just a privileged few. The organization's initial mission was to create and disseminate advanced AI technologies in a way that prioritized public benefit over profit. OpenAI aimed to foster collaboration and transparency, making its cutting-edge research and models available for anyone to use, modify, and build upon. This philosophy was deeply rooted in the belief that AI should be a tool for collective advancement, addressing pressing global challenges such as climate change, healthcare, education, and economic inequality. OpenAI's early emphasis on open-source aligned with the broader movement of democratizing knowledge and technology, aiming to prevent the concentration of power in the hands of a few entities, be they governments or corporations.

Empowering Local Communities and Individuals through Open-Source AI

The open-source nature of AI held the potential to profoundly transform how local communities and individuals addressed challenges, particularly in the realm of civil defense. OpenAI's initial vision suggested that by making powerful AI models openly available, communities would be able to deploy AI solutions locally, independent of centralized systems or proprietary platforms. This capability would be invaluable during times of crisis, when infrastructure may be compromised, and communities need to rely on decentralized tools to manage critical resources, coordinate disaster response, and even save lives.

In civil defense scenarios—whether responding to natural disasters, pandemics, or conflicts—open-source AI could empower local governments, grassroots organizations, and even individual citizens with powerful tools for real-time decision-making. For example, AI models could assist in disaster prediction, early warning systems, emergency medical triage, and the allocation of limited resources like food, water, and shelter. By allowing anyone with access to adequate hardware to run these models, communities could become more resilient and self-sufficient, able to respond to crises even when cut off from broader support systems or centralized authorities.

Moreover, open-source AI would enable equitable access. In a world where technological advancements are often concentrated in wealthier, more developed regions, an open AI platform would help level the playing field. Communities that might otherwise be left behind due to financial or infrastructure constraints could now harness the same powerful AI tools as more affluent nations. This democratization of AI could be particularly important for small or under-resourced communities that might not have the capital to afford proprietary AI services or cloud-based models. Through open access, they would have the autonomy to deploy and adapt AI tools according to their specific needs, whether for disaster preparedness, conflict resolution, or community development.

Historical Successes of Open-Source Software in Strengthening Communities

OpenAI's initial open-source vision was not without precedent. The open-source movement has historically played a transformative role in the technology landscape, providing powerful examples of how freely available software can

strengthen communities, promote innovation, and foster collaboration on a global scale.

One of the most iconic successes of open-source software is the Linux operating system. Created by Linus Torvalds in the early 1990s and released as free and open-source software, Linux quickly gained traction and became the backbone of much of the world's computing infrastructure. Its open-source nature allowed developers from around the globe to contribute to its development, resulting in a robust, secure, and flexible platform that powers everything from personal computers to enterprise servers and even the majority of the world's cloud services. The success of Linux demonstrated the power of community-driven development, where the collective input of many individuals and organizations led to a software platform that rivaled and often surpassed proprietary alternatives.

Similarly, open-source cryptography has been crucial in securing digital communications globally. Open-source cryptographic libraries such as OpenSSL have provided the tools necessary to secure internet traffic, protect sensitive data, and enable privacy for individuals and organizations. The openness of these tools allows experts from around the world to audit, improve, and adapt them for various applications, ensuring that everyone has access to strong security protocols, not just large corporations or government entities. This widespread access to secure cryptographic tools has been essential for the development of secure online banking, e-commerce, and digital communications.

These historical examples show how open-source software, by being freely accessible and modifiable, empowers communities to take control of their own technological infrastructure. Whether it's enabling secure communications in conflict zones or providing the operating systems that power disaster response systems, open-source software has consistently demonstrated its capacity to foster resilience, innovation, and equitable access to technology.

Had OpenAI remained committed to its open-source mission, its AI models could have followed the same path, becoming vital tools in civil defense and beyond. Local communities and individuals could have harnessed AI to address the specific challenges they face, without relying on expensive or inaccessible proprietary services. Instead, OpenAI's shift away from open-source has created significant

barriers to realizing this vision, leaving much of its transformative potential untapped, particularly for those who stand to benefit the most.

In summary, OpenAI's original vision of open-source AI held the promise of empowering communities around the world by providing equitable access to powerful tools that could enhance local resilience, self-sufficiency, and crisis preparedness. The success of past open-source initiatives like Linux and open-source cryptography demonstrates how a commitment to openness and accessibility can strengthen communities, encourage innovation, and foster widespread collaboration. OpenAI had the potential to do the same with AI—but that promise remains unfulfilled as the organization has shifted its focus toward profit-driven models.

The Shift to Profit-Driven Models

OpenAI's journey from its founding as a non-profit to its eventual transition into a capped-profit organization represents a significant turning point in its mission and approach. When OpenAI was first established, it was a non-profit organization with a clear focus on ensuring that artificial intelligence would be developed for the benefit of all humanity, without prioritizing financial returns. However, in 2019, OpenAI underwent a structural transformation, transitioning to a new model called a "capped-profit" organization. This shift allowed OpenAI to accept outside investments, promising investors a capped return of up to 100 times their investment while theoretically maintaining a long-term focus on the public good. Although the capped-profit model was designed to balance financial incentives with OpenAI's broader mission, in practice, it marked a significant departure from the organization's original commitment to openness and accessibility.

Transition from Non-Profit to Capped-Profit: A Strategic Pivot

OpenAI's decision to transition to a capped-profit model was driven largely by the immense costs associated with developing cutting-edge artificial intelligence. Building large-scale models like GPT-4 requires not only a team of world-class researchers but also vast computational resources. As the scale of OpenAI's ambitions grew, so did its need for capital, which exceeded the limits of typical non-profit fundraising mechanisms. By allowing investors to participate through

the capped-profit structure, OpenAI gained access to billions of dollars in capital, including a highly publicized $1 billion investment from Microsoft.

While this injection of funds enabled OpenAI to accelerate its research and develop more sophisticated models, it also introduced a new set of financial incentives that ultimately shifted the organization's focus. The need to provide returns to investors placed greater emphasis on monetizing AI technologies, steering OpenAI away from its initial commitment to open-source and free distribution. The resulting shift from a public-benefit model to one that also had to satisfy investor interests altered the organization's trajectory and the accessibility of its models.

Monetization Strategies: API Pricing and the Move Away from Open Distribution

As OpenAI continued to advance its technology, it introduced monetization strategies that limited access to its most powerful models. One of the primary ways OpenAI began monetizing its AI was through the introduction of paid API access. Rather than distributing the models openly, as originally promised, OpenAI began offering access to GPT models via a cloud-based API service. Users could integrate GPT-3, GPT-4, and other models into their applications by paying for usage through a subscription model or usage-based pricing.

This shift had profound implications for the accessibility of OpenAI's models. Under the paid API model, only individuals, organizations, or businesses with the financial means to afford the service could access OpenAI's cutting-edge technology. While OpenAI still allowed some limited free access, the full potential of GPT models became effectively locked behind a paywall. The more advanced and resource-intensive the AI model, the more expensive it became to use. As a result, this approach significantly restricted access for smaller organizations, nonprofits, educational institutions, and individuals who lacked the financial resources to pay for API services.

In addition to API pricing, OpenAI imposed licensing restrictions on its models. For example, GPT-4 was not made available for direct download or local deployment; it could only be accessed via the cloud-based API. This decision further constrained the ability of communities, particularly those in crisis or low-resource settings, to deploy AI models locally. Even though devices like the Seeed Orin and other powerful hardware were becoming affordable enough to run large AI

models locally, OpenAI's decision to maintain control over distribution via cloud services meant that communities could not access these models without an internet connection and the financial means to pay for API usage. This further undermined the original vision of open-source AI for the public good.

How These Decisions Have Limited Access to GPT-4

OpenAI's decision to prioritize API pricing and cloud-based access has created significant barriers for deploying GPT-4 at the local level, particularly in the context of civil defense and crisis preparedness. The need for paid access to the models creates a cost barrier for individuals, grassroots organizations, and communities that could benefit from deploying AI in emergency situations. In scenarios where internet access is compromised—such as during natural disasters or conflicts—the cloud-based nature of OpenAI's services further limits the practical utility of GPT-4. Communities that may have the hardware to run these models locally are nonetheless unable to do so because OpenAI does not offer local deployment options for its most advanced models.

Additionally, this move away from open-source distribution stifles the kind of collaborative innovation that open-source software historically enables. Under the current model, developers cannot download or modify GPT-4 to adapt it for specific use cases, such as improving its performance in low-resource settings, tailoring it for particular languages or dialects, or refining it for specialized tasks like disaster response or medical triage. The centralized control of the models limits the opportunity for community-driven improvements that could have been possible in an open-source framework.

In the context of civil defense, these limitations have profound consequences. OpenAI's profit-driven model restricts access to GPT-4 for local communities that may otherwise be able to use it to enhance emergency response capabilities, resource management, or even predictive modeling for disaster scenarios. Moreover, it exacerbates existing inequalities, where wealthier communities or organizations with the financial means to pay for cloud access are able to leverage GPT-4's capabilities, while those without the resources are left behind.

By moving away from its original vision of open-source AI, OpenAI has placed its most powerful tools out of reach for many, undermining the potential of GPT-4 to serve as a critical resource for civil defense. While the capped-profit model may

have enabled OpenAI to attract the capital it needed to continue developing cutting-edge AI, it has done so at the cost of accessibility, collaboration, and equity, limiting the ability of individuals and communities to deploy AI in ways that could truly benefit society.

In summary, OpenAI's shift from a non-profit organization to a capped-profit entity has driven the introduction of monetization strategies, such as API pricing and licensing restrictions, that limit access to its most powerful models, including GPT-4. These decisions have created barriers to local deployment, especially for individuals and communities in need of decentralized, resilient AI solutions. As a result, the promise of open-source AI has been compromised, and the transformative potential of these models, particularly in civil defense, remains unrealized.

How OpenAI's Greed Is Preventing Local Deployment

As OpenAI has evolved from its original non-profit mission to a capped-profit organization, its financial priorities have increasingly driven decisions about how its AI models, such as GPT-4, are accessed and used. While the technology behind GPT-4 could be a game-changer for civil defense and local deployment in times of crisis, OpenAI's licensing costs and restrictive access policies have effectively prevented individuals, communities, and even smaller organizations from harnessing its full potential. This shift from an open-source ethos to a profit-driven model has deprived the public of valuable tools that could be used to save lives and bolster community resilience in the face of disasters and emergencies.

Licensing Costs and Restrictions Imposed by OpenAI on Accessing and Using GPT-4

OpenAI's decision to impose strict licensing and access restrictions on its models has become one of the primary obstacles to local deployment. Unlike earlier AI models that could be downloaded and run locally on available hardware, GPT-4 is accessible only via OpenAI's cloud-based API. This API access is controlled through a tiered pricing structure, with free or limited usage tiers offering minimal functionality, while the most powerful capabilities are locked behind paid subscriptions and usage fees.

These licensing costs can be prohibitively expensive for many users. Individuals, nonprofits, and small organizations—particularly those operating in low-resource environments—are often unable to afford the ongoing subscription fees required to access GPT-4. For example, civil defense organizations that could benefit from GPT-4's ability to assist in real-time crisis management, medical advice, and logistics would need to allocate significant portions of their budgets to cover the costs of API access. This creates a barrier to entry that restricts access to a powerful tool at a time when it is most needed.

Furthermore, OpenAI's refusal to allow local deployment of GPT-4 exacerbates this problem. Even for those who could afford the necessary hardware—such as devices like the Seeed Orin, which are capable of running AI models locally—OpenAI's licensing restrictions prevent users from downloading and running GPT-4 on their own systems. This means that all interactions with GPT-4 must go through OpenAI's servers, which not only incurs ongoing financial costs but also limits flexibility, control, and accessibility, particularly in environments where reliable internet access is unavailable.

In a time of crisis, where infrastructure might be damaged or disrupted, reliance on cloud-based services creates an additional point of failure. Communities or emergency response teams that cannot maintain a stable connection to OpenAI's servers are left without access to the tools they need to coordinate disaster response, manage resources, or offer critical medical assistance. OpenAI's decision to prioritize centralized, cloud-based control over the distribution of its models directly undermines the ability of users to deploy AI locally in the very situations where it could make the most difference.

Financial Incentives Driving the Shift Toward Cloud-Based Access Instead of Enabling Local Use

The financial incentives driving OpenAI's cloud-based business model are rooted in the company's shift toward a profit-driven approach. By centralizing access to its AI models through a paid API, OpenAI can generate consistent revenue streams from users who are dependent on its servers for their AI needs. This subscription-based model is a departure from the organization's original open-source philosophy, and it prioritizes the financial interests of the organization and its investors over the broader public good.

Cloud-based services offer several financial advantages to OpenAI. First, they allow the organization to maintain control over its intellectual property, ensuring that users pay for access and preventing them from reverse-engineering or modifying the models. Second, a subscription model generates ongoing revenue, rather than a one-time purchase fee, creating a more sustainable and lucrative business model. Third, centralized control allows OpenAI to gather data on how its models are used, which can further drive product improvements and marketing strategies.

However, these financial incentives come at the cost of accessibility and flexibility for users. Local use of GPT-4, which would allow individuals and organizations to download and run the model on their own hardware, is not just a technical possibility but also a much-needed capability for many sectors, including civil defense. Local deployment would offer critical advantages in terms of independence from internet connectivity, security, and the ability to modify the model for specific use cases. By maintaining exclusive control over the cloud infrastructure, OpenAI is prioritizing its revenue streams over the potential life-saving benefits of locally deployed AI.

This cloud-first approach also limits innovation. In an open-source environment, users could customize AI models to better suit local needs, optimize them for specific hardware configurations, or integrate them with other systems in ways that a cloud-based service cannot accommodate. For example, in rural areas with limited connectivity, local use of GPT-4 could be integrated into offline emergency systems, empowering communities to use AI for crisis prediction, disaster recovery, or even medical diagnostics without needing to rely on an unstable internet connection. But the financial motivations behind OpenAI's decisions prevent such use cases from being realized.

Missed Opportunities for Local Civil Defense Due to Lack of Open Access

The consequences of OpenAI's cloud-based, profit-driven model are especially apparent when considering the missed opportunities for local civil defense. Civil defense organizations, especially those operating in disaster-prone or conflict regions, often have limited access to real-time information, resources, and decision-making tools. The deployment of AI at a local level could dramatically enhance their capacity to respond to crises. Yet, without the ability to download

and use GPT-4 locally, these organizations are left with fewer options and greater vulnerabilities.

1. Natural Disaster Preparedness and Response

In the event of a natural disaster such as an earthquake, hurricane, or wildfire, access to real-time information can be the difference between life and death. AI models like GPT-4 could assist in crisis prediction, helping emergency teams anticipate events, allocate resources, and make informed decisions under pressure. With open access and local deployment, GPT-4 could be used to process sensor data from local weather stations, monitor communications, and provide predictive modeling to help communities better prepare for, and respond to, natural disasters.

However, under OpenAI's current cloud-based model, such real-time functionality is hindered by the need for continuous connectivity and ongoing subscription fees. In areas where communication infrastructure is compromised, reliance on cloud services is unreliable. Communities are thus deprived of a tool that could have been available locally, offering them the ability to process data and run simulations even in the absence of internet access.

2. Emergency Medical Advice and Triage

Another critical application of AI in civil defense is in the field of emergency medicine. AI models like GPT-4 can offer real-time medical advice, especially in regions where healthcare resources are strained or nonexistent. In a disaster scenario, where medical professionals may be overwhelmed or unable to reach affected areas, GPT-4 could help guide non-expert first responders in performing triage, administering basic care, and stabilizing patients until professional help arrives.

In communities with limited internet access or unstable connectivity, local deployment of AI could be lifesaving. GPT-4, running on local hardware, could be used to answer medical questions, guide emergency procedures, and offer psychological support to trauma victims. But because OpenAI's model is cloud-based and requires paid access, many communities that could benefit from this technology are left without these critical capabilities.

3. Resource Allocation and Crisis Management

Resource allocation is another area where AI could significantly enhance civil defense. In crisis scenarios, efficient distribution of resources such as food, water, shelter, and medical supplies is crucial. AI models can analyze logistical data, prioritize resource distribution, and assist in optimizing supply chains to ensure that aid reaches the most vulnerable populations as quickly as possible. Local AI deployment would allow communities to manage these processes even when disconnected from central command structures or the internet.

However, the financial model adopted by OpenAI has effectively locked this capability behind a paywall, preventing many communities from deploying GPT-4 for resource management when it is most needed.

Conclusion

In conclusion, OpenAI's shift toward profit-driven models, reflected in restrictive licensing, high API costs, and the lack of local deployment options, has significantly hindered the widespread use of GPT-4 for civil defense. Financial incentives have driven the shift toward centralized cloud-based access, prioritizing revenue generation over the potential life-saving benefits of local AI deployment. As a result, countless communities are missing out on the opportunity to use AI for disaster preparedness, emergency medical care, and resource management, all of which could significantly enhance resilience in times of crisis. The missed opportunities are a stark reminder that the pursuit of profit has come at the expense of public good, particularly in sectors like civil defense where access to cutting-edge AI tools could make a critical difference.

The Consequences for Civil Defense

Artificial intelligence (AI) has the potential to revolutionize civil defense, enabling communities to become more resilient and self-sufficient in the face of crises such as natural disasters, pandemics, and conflicts. By decentralizing decision-making and providing advanced tools for real-time response, AI could empower local communities, emergency responders, and civil defense organizations to act swiftly and effectively during emergencies. However, OpenAI's shift toward a profit-driven model, which restricts access to its most advanced AI systems like

GPT-4, has severely limited the ability of communities to harness these technologies. The consequences of this shift are profound, especially when considering how AI could otherwise enhance local civil defense efforts, improve crisis management, and save lives.

The Critical Role AI Could Play in Decentralizing and Strengthening Civil Defense Efforts

One of the most powerful aspects of AI is its ability to process vast amounts of information in real time, identify patterns, and make data-driven decisions quickly and accurately. In civil defense, where time and accuracy are often the most crucial factors in saving lives, AI could be an invaluable asset. By decentralizing the deployment of AI tools—allowing communities to access and use these models locally—civil defense efforts could be significantly enhanced.

Local deployment of AI models would allow communities to operate independently, without needing to rely on centralized systems or external networks that may be unavailable during crises. For example, in the immediate aftermath of a natural disaster or during a cyberattack, communication infrastructure may be damaged, and access to the internet or centralized cloud services could be interrupted. In such situations, having locally available AI models, running on hardware like the Seeed Orin, could provide crucial support for decision-making and resource management, even when communities are isolated or cut off from external aid.

Decentralized AI models could enable more efficient coordination among local authorities, emergency services, and volunteer organizations. Rather than waiting for instructions or resources from centralized government agencies, communities could use AI to assess the situation on the ground, prioritize actions, and distribute resources where they are needed most. This level of autonomy and real-time responsiveness could make the difference between an organized, effective response and a chaotic, delayed one.

Examples of How Local AI Models Could Assist in Crisis Situations

The use of AI in civil defense spans a wide range of applications, each of which could benefit from the local deployment of advanced models like GPT-4. Some of

the most critical areas where AI could play a role include resource management, medical triage, and real-time decision-making during emergencies.

1. **Resource Management**
    - In any crisis, the efficient allocation of resources is paramount. Whether it is distributing food and water after a natural disaster, coordinating shelter during a mass evacuation, or ensuring that medical supplies reach the most vulnerable populations, AI can help optimize the logistics of resource management.
    - AI models could analyze data from various sources—such as geographic information systems (GIS), satellite imagery, social media reports, and sensor networks—to determine where resources are most urgently needed. These models could then recommend the best routes for delivering supplies, identify bottlenecks in transportation, and help prioritize distribution based on real-time information.
    - By running these AI models locally, communities could avoid the delays and inefficiencies that often come with waiting for centralized agencies to coordinate responses. The ability to process information quickly and act independently would be especially important in remote or underserved areas, where external support may be slow to arrive.
2. **Medical Triage**
    - In the chaotic aftermath of a disaster or during a large-scale medical emergency, such as a pandemic, emergency responders are often overwhelmed by the sheer number of patients needing care. AI could assist in medical triage, helping to prioritize patients based on the severity of their injuries or symptoms and ensuring that limited medical resources are used most effectively.
    - A locally deployed AI model like GPT-4 could be used to analyze patient data, suggest diagnoses, and recommend treatment protocols. For example, AI could assist paramedics in assessing which patients need immediate attention, which can wait, and which could

be treated on-site without requiring hospitalization. This would reduce the burden on hospitals and medical centers, which are often inundated during crises.

- In regions where healthcare professionals are scarce, AI could also guide non-expert first responders in performing critical medical interventions, offering real-time advice on administering first aid, stabilizing patients, or performing life-saving procedures. By using AI locally, communities could provide emergency medical care even when they are cut off from centralized healthcare systems.

3. **Real-Time Decision-Making**

    - During an unfolding crisis, decision-makers often face a flood of information, some of which may be conflicting or incomplete. AI models can help sift through this information, identify the most relevant data, and provide recommendations for action based on real-time analysis.
    - For example, in the event of a natural disaster, AI could analyze weather data, damage reports, and communication patterns to help authorities decide where to deploy emergency personnel, how to organize evacuations, and which infrastructure to prioritize for repairs. In conflict zones, AI could assist in tracking the movement of displaced populations, identifying safe zones, and coordinating humanitarian aid.
    - With local AI deployment, these real-time decision-making tools could be used even when communities are disconnected from external command centers. This level of autonomy could empower local leaders to make informed, data-driven decisions without having to wait for instructions from higher authorities, potentially saving lives in the critical early hours of a disaster.

How OpenAI's Profit-Driven Approach Hinders Local Deployment and Crisis Response

Despite the clear potential of AI to enhance civil defense, OpenAI's shift toward a profit-driven approach has hindered the ability of communities to deploy these

models locally. As previously discussed, OpenAI's business model revolves around cloud-based access to its AI models through a paid API service, which requires ongoing subscriptions and internet connectivity. This creates several key barriers to local deployment, particularly in crisis scenarios where reliance on external servers and internet access may be impractical or impossible.

1. **Dependence on External Servers**
    - By restricting access to GPT-4 through a cloud-based API, OpenAI has made it difficult for communities to use these models independently. In crisis situations, such as natural disasters or cyberattacks, internet connectivity may be disrupted, making it impossible to access cloud-based AI services.
    - Even in cases where the internet is available, latency and server outages can make real-time decision-making difficult. For example, in a rapidly evolving disaster scenario, waiting for data to be processed by remote servers can introduce delays that could be fatal. Local deployment of AI models, which could run directly on community hardware, would avoid these issues, allowing for uninterrupted use of AI tools during emergencies.
2. **High Costs of Cloud-Based Access**
    - OpenAI's pricing structure for accessing GPT-4 through its API is prohibitive for many smaller organizations, nonprofits, and local governments that could benefit most from AI in civil defense. The cost of ongoing API usage can quickly add up, especially for organizations with limited budgets. As a result, many communities are effectively priced out of using this technology, despite its potential to save lives and improve crisis management.
    - Local deployment, by contrast, would involve a one-time investment in hardware capable of running AI models, such as the Seeed Orin, rather than ongoing subscription fees. This would make AI more accessible to communities with limited financial resources, allowing them to deploy AI tools without worrying about the long-term costs of cloud services.

3. **Lack of Flexibility and Control**
    - OpenAI's cloud-based model also limits the flexibility of how AI can be used in different contexts. Local deployment would allow communities to customize AI models to suit their specific needs, whether that involves adapting the model for particular languages, integrating it with local sensor networks, or modifying it for use in specific types of disasters.
    - By maintaining control over GPT-4's deployment through its servers, OpenAI prevents communities from modifying the model to fit their local circumstances. This lack of control limits the potential for innovation and reduces the overall utility of AI in civil defense, as users are forced to rely on a one-size-fits-all solution that may not be optimized for their unique needs.

Conclusion

OpenAI's shift toward a profit-driven model, with a focus on cloud-based access and subscription pricing, has created significant barriers to the local deployment of AI models like GPT-4, particularly in the context of civil defense. By preventing communities from accessing and using AI locally, OpenAI's business model hinders the ability of communities to prepare for and respond to crises independently, especially in situations where reliance on external servers and internet access is impractical. The consequences of this approach are clear: missed opportunities to enhance resource management, improve medical triage, and enable real-time decision-making in emergencies. If AI is to fulfill its potential as a tool for civil defense, OpenAI must reconsider its profit-driven approach and explore ways to make its models more accessible to local communities.

A Return to Open-Source: The Ethical Imperative

Artificial intelligence (AI) holds immense potential for improving global resilience, especially in the face of crises such as natural disasters, public health emergencies, and conflicts. However, the current profit-driven business model embraced by organizations like OpenAI has created significant barriers to accessing these powerful tools, particularly for communities that are most

vulnerable to such crises. A return to an open-source model for AI, particularly for models like GPT-4, is not only an ethical imperative but also a practical solution for ensuring that AI fulfills its potential as a public good. In the context of civil defense, open-source AI could democratize access, allowing communities to enhance their preparedness and resilience in ways that are currently inaccessible due to financial and technical barriers.

The Ethical Argument for Making AI Like GPT-4 Open-Source in the Context of Civil Defense

The ethical case for open-source AI, particularly for civil defense, is grounded in the idea that life-saving technology should not be restricted to those with the financial means to afford it. AI models like GPT-4 have the potential to provide critical assistance during crises, yet their deployment is currently limited to those who can pay for cloud-based access and manage the ongoing costs associated with API usage. This creates a significant disparity between wealthier organizations and communities, which can afford to utilize advanced AI tools, and poorer or more isolated communities, which are left without access to these life-saving technologies. This inequality is not just a technological issue but a moral one, as it perpetuates existing disparities in global resilience to disasters and emergencies.

The ethical foundation of open-source AI lies in the principle of equitable access to technology that can serve the common good. When AI is developed with a profit-driven motive, as in the case of OpenAI's current business model, it becomes a commodity accessible only to those who can afford it. However, AI—particularly in contexts such as civil defense, where lives are at stake—should be treated as a public good, accessible to all regardless of their financial means. The ability to deploy AI locally, without being dependent on expensive cloud services or ongoing subscription fees, is a vital component of this argument. Open-source AI would allow communities to take ownership of their own civil defense strategies, tailoring AI tools to meet their specific needs and reducing their dependence on external entities for critical technologies.

Moreover, open-source AI would ensure that communities and individuals have the freedom to audit, improve, and modify AI models to suit their unique requirements. This transparency is crucial for ensuring that AI is used ethically and

responsibly. In a closed, cloud-based model, the algorithms and decision-making processes used by AI models are opaque to most users, which can lead to mistrust, misapplication, and, in some cases, unethical use. By making AI models open-source, the community at large—including researchers, developers, and civil defense experts—can collaborate to improve the models, making them more robust, reliable, and responsive to the ethical concerns that arise in different contexts.

How Open-Source AI Could Democratize Access and Allow Communities to Build Local Resilience

A return to open-source AI would democratize access, allowing communities to build local resilience by deploying, customizing, and using AI models independently. Currently, the centralized control that OpenAI exerts over models like GPT-4 restricts how and where they can be deployed, effectively locking out many of the communities that could benefit the most. By contrast, open-source AI would enable local deployment on affordable hardware, allowing communities to operate independently of external networks and cloud services, which are often unavailable or unreliable in the midst of a crisis.

This shift toward open-source AI would also foster innovation at the local level. Communities, developers, and small organizations could modify open-source models to better address their specific needs. For example, a community in a flood-prone area could modify an AI model to integrate real-time data from local weather stations and river gauges, allowing for highly localized flood predictions. In a remote village without access to constant healthcare, AI models could be adapted to provide medical diagnostics or triage guidance based on locally available resources and health data. Such modifications are not possible under the current cloud-based model, where AI tools are designed for broad, generic applications and cannot be customized to fit unique local conditions.

Furthermore, open-source AI would empower communities to operate independently during crises. By downloading and running AI models on local hardware, communities could continue to use AI for resource management, medical triage, and decision-making even when they are cut off from the internet or external servers. This ability to operate in “offline mode” is critical during natural disasters, where infrastructure may be damaged, or during cyberattacks,

where communication systems are deliberately disrupted. Local resilience would be significantly enhanced if communities had access to AI tools that were not contingent on external connectivity.

Examples of Industries and Initiatives that Thrive Due to Open-Source Models

The potential for open-source AI to transform civil defense is not merely speculative; there are numerous examples of industries and initiatives that have thrived due to the adoption of open-source models. These examples demonstrate the power of open access to technology in fostering innovation, collaboration, and resilience across various fields.

1. **Healthcare**
    - Open-source initiatives in healthcare have made significant contributions to improving medical care, particularly in low-resource settings. For example, OpenMRS, an open-source medical record system, has been widely adopted in developing countries, where proprietary solutions were either too expensive or ill-suited to local healthcare systems. OpenMRS allows healthcare providers to customize the platform to meet their needs, enabling more efficient patient tracking, data analysis, and resource management in settings where these capabilities were previously lacking.
    - Open-source medical tools have also contributed to advancements in diagnostics, disease tracking, and healthcare delivery. During the COVID-19 pandemic, open-source models for tracking virus transmission and predicting hospital capacity helped communities around the world respond more effectively to the crisis. These initiatives highlight the power of open-source technology in delivering life-saving solutions that are accessible to all, regardless of financial resources.
2. **Cybersecurity**
    - The cybersecurity industry has long relied on open-source tools to protect critical systems from cyber threats. Tools such as OpenSSL (used to secure internet communications) and Wireshark (a network analysis tool) have been developed collaboratively by the global

community and are freely available to anyone. Open-source cybersecurity tools allow users to audit and improve security protocols, ensuring that they can be trusted and are continually updated to respond to new threats.

- The success of open-source cybersecurity tools demonstrates the importance of transparency and collaboration in fields where trust is essential. Just as communities around the world rely on open-source tools to protect their networks from cyberattacks, they could similarly benefit from open-source AI tools to enhance their civil defense capabilities.

3. **Education**

    - The open-source movement has also played a critical role in democratizing education. Platforms like Moodle, an open-source learning management system, have made it possible for educational institutions around the world to provide high-quality online education at minimal cost. By allowing institutions to customize the platform, integrate local content, and adapt the system to specific teaching needs, open-source educational tools have expanded access to education in places where proprietary systems were either too costly or unavailable.
    - In the context of civil defense, open-source AI could similarly provide tools for education and training in disaster preparedness. Communities could use AI to simulate crisis scenarios, train emergency responders, and educate the public about how to respond in emergencies—all without the financial or technological barriers imposed by proprietary systems.

These examples from healthcare, cybersecurity, and education demonstrate the transformative potential of open-source models. By making advanced technology freely accessible, open-source initiatives foster collaboration, innovation, and resilience. A return to open-source AI for civil defense would build on these successes, ensuring that communities around the world can access the tools they need to protect themselves, respond to crises, and build a safer future.

Conclusion

The ethical case for open-source AI is clear: in a world where crises are becoming more frequent and severe, it is morally unjustifiable to restrict access to life-saving technology based on financial means. Open-source AI could democratize access, empowering communities to build local resilience, innovate solutions to their unique challenges, and operate independently during emergencies. The success of open-source initiatives in healthcare, cybersecurity, and education demonstrates that this model not only fosters innovation but also ensures that technology serves the public good, rather than being controlled by a small group of corporations. If OpenAI were to return to its original open-source vision, it could help build a more resilient, equitable world, where AI serves all of humanity, not just those who can afford it.

Solutions to Overcome the Profit Barrier

The shift by OpenAI toward profit-driven models has created a significant barrier to the local deployment of AI, particularly for civil defense applications where advanced tools like GPT-4 could save lives and enhance resilience. However, there are solutions to overcome this profit barrier, including exploring open-source alternatives, increasing public pressure and advocacy for OpenAI to return to its founding principles, and seeking government or non-profit funding to support the development of freely available AI. By focusing on these solutions, we can mitigate the effects of OpenAI's profit-driven approach and ensure that communities, particularly those facing crises, have access to powerful AI tools without being financially burdened.

1. Open-Source Alternatives to OpenAI That Could Provide Similar Capabilities

One of the most promising solutions to the profit barrier imposed by OpenAI is the development and use of open-source AI alternatives. While OpenAI's GPT-4 remains one of the most advanced language models, the open-source AI community has made significant strides in creating comparable models that are freely available to the public.

a. GPT-J and GPT-Neo

- GPT-J and GPT-Neo, developed by EleutherAI, are two open-source alternatives that provide functionality similar to OpenAI's GPT-3 and GPT-4 models. EleutherAI is a collective of researchers dedicated to creating open and unrestricted access to large-scale language models. GPT-J and GPT-Neo are both available for free, and unlike OpenAI's models, they can be downloaded and run on local hardware without the need for ongoing subscription fees or cloud-based services.
- These models may not match GPT-4 in terms of size and complexity, but they are powerful enough for many applications, including civil defense. Communities can deploy these models locally for tasks such as resource management, crisis communication, and emergency decision-making. Moreover, because these models are open-source, users can modify and adapt them to better suit local needs, ensuring that the models are flexible and responsive to the unique challenges faced by different communities.

b. BLOOM

- BLOOM is another open-source language model developed by the BigScience project, a collaborative effort involving researchers from around the world. BLOOM is designed to provide large-scale natural language processing capabilities similar to GPT-3 and GPT-4 but with a focus on being multilingual and inclusive of a wide range of languages and dialects.
- Like GPT-J and GPT-Neo, BLOOM is open-source and freely available, meaning that communities can download and run the model on local hardware. This is particularly important for civil defense efforts in non-English-speaking regions, where local dialects and languages are critical for effective communication during crises. BLOOM's open-source nature ensures that the model can be fine-tuned and adapted to suit specific linguistic and cultural contexts, making it a valuable tool for enhancing global resilience.

c. OpenAI Competitors and Ecosystem

- Other organizations and research collectives are actively working on creating open-source AI models that rival OpenAI’s. Hugging Face, for

example, hosts a wide range of open-source AI models that can be downloaded, modified, and deployed locally. Hugging Face has established itself as a key player in the open-source AI ecosystem, providing tools and models that can be easily integrated into local systems.

- OpenAI's focus on centralized cloud-based services opens the door for these competitors to develop decentralized AI solutions, which offer both technical and ethical advantages by prioritizing accessibility and public good over profit. As these models continue to improve, they will become increasingly viable alternatives for communities seeking to deploy AI for civil defense without being tied to OpenAI's profit-driven framework.

2. Public Pressure and Advocacy for OpenAI to Return to Its Founding Principles

While open-source alternatives provide a practical solution to the profit barrier, another important strategy is to advocate for OpenAI to return to its original mission of making AI accessible for the public good. Public pressure from various sectors—including the research community, civil society, and policymakers—could push OpenAI to reconsider its current approach and prioritize broader societal benefit over profit maximization.

a. Grassroots Advocacy

- Grassroots advocacy can play a powerful role in shaping public opinion and influencing corporate behavior. Researchers, activists, and civil defense organizations could collaborate to raise awareness about the barriers created by OpenAI's profit-driven model and the missed opportunities for using AI in life-saving applications. Public campaigns highlighting these issues could pressure OpenAI to reconsider its licensing policies, pricing structures, and cloud-based access requirements, advocating for more equitable access to advanced AI models.
- Social media and public platforms can be used to amplify these concerns, drawing attention to the ethical implications of restricting access to AI technologies that could serve the common good. If enough voices call for OpenAI to re-align with its original open-source principles, it may lead to changes in its business model or the creation of new policies aimed at balancing profit with public interest.

b. Policy Advocacy and Government Regulation

- Another avenue for public pressure involves advocating for government policies that regulate the accessibility of advanced AI models, particularly in sectors like civil defense, healthcare, and education. Governments could introduce policies that require companies like OpenAI to make critical technologies available to the public, either through open-source distribution or through subsidized access for specific applications like emergency preparedness.
- These policies could incentivize companies to adopt more equitable distribution models, ensuring that AI technologies are available to underserved communities and organizations that lack the financial resources to pay for subscription-based services. Governments could also establish guidelines to ensure that essential AI tools, particularly those with life-saving potential, are treated as public goods rather than commercial products.

c. Collaboration with Non-Profits and Ethical AI Organizations

- Ethical AI organizations such as the Partnership on AI and the Algorithmic Justice League have long advocated for the responsible and equitable development of AI. Collaboration with these organizations could amplify the call for OpenAI to embrace its original open-source principles and provide more accessible AI tools.
- By engaging with these advocacy groups, researchers and civil defense organizations can push for systemic changes in the way AI is developed and distributed, ensuring that advanced AI models like GPT-4 are made available for public use without the financial barriers imposed by a profit-driven approach.

3. Potential for Government or Non-Profit Funding to Support the Development of Freely Available AI for Civil Defense

Governments and non-profit organizations can play a crucial role in overcoming the profit barrier by funding the development of open-source AI models specifically tailored for civil defense applications. Such funding could create a new

class of AI tools that are freely available to the public, ensuring that communities around the world can access and deploy AI in times of crisis.

a. Government-Led AI Initiatives

- Governments have the capacity to fund large-scale AI projects that prioritize public good over profit. By allocating resources toward the development of open-source AI models for civil defense, governments could ensure that their populations have access to the tools necessary to respond effectively to disasters, emergencies, and other crises.
- These government-led AI initiatives could focus on developing models that are optimized for local deployment, particularly in regions with limited infrastructure or connectivity. Governments could also create open repositories where AI models are made available to the public for free, much like open data initiatives that have provided valuable resources for researchers, innovators, and emergency response organizations.
- Additionally, governments could establish grants and incentives for researchers and developers working on open-source AI projects, fostering an ecosystem where innovation is driven by the desire to serve the public good rather than the pursuit of profit.

b. Non-Profit Funding and Philanthropic Support

- Non-profit organizations and philanthropic foundations can also play a significant role in funding the development of freely available AI models. Organizations like the Bill & Melinda Gates Foundation, which has a history of funding public health and education initiatives, could invest in the creation of AI tools designed for civil defense.
- Non-profit funding could enable the development of specialized AI models that are tailored to the needs of underserved or vulnerable populations. For example, a non-profit could fund the development of AI tools that assist in resource management during natural disasters, help provide medical triage in conflict zones, or offer psychological support during crises. These models could be distributed freely to communities around the world, empowering them to take ownership of their civil defense strategies.

- Philanthropic support could also help cover the costs of maintaining and updating open-source AI models, ensuring that they remain relevant and effective as new challenges and technologies emerge.

c. International Collaboration for AI Development

- International organizations such as the United Nations or the World Health Organization could spearhead collaborative efforts to fund and develop open-source AI models for civil defense. By pooling resources from governments, non-profits, and private sector entities, these organizations could create globally accessible AI tools that are designed to enhance resilience and disaster preparedness in vulnerable regions.
- International collaboration would also ensure that AI models are developed with a global perspective in mind, accounting for the diverse needs of different cultures, languages, and regions. This would help ensure that the models are inclusive and can be adapted to serve a wide range of civil defense applications across different contexts.

Conclusion

Overcoming the profit barrier that currently limits access to AI tools like GPT-4 is not only possible but essential for ensuring that communities around the world can benefit from the life-saving potential of AI. By embracing open-source alternatives, increasing public pressure and advocacy for equitable access to AI, and securing government or non-profit funding for the development of freely available models, we can create a future where AI serves the public good rather than being restricted by profit motives. These solutions offer a path forward that prioritizes resilience, innovation, and the ethical use of AI for civil defense and beyond.

Ethical and Strategic Considerations

As artificial intelligence (AI) continues to advance and its applications spread across critical sectors such as healthcare, civil defense, and disaster management, the ethical and strategic considerations surrounding its development and deployment become increasingly important. The rise of profit-driven models, like

those adopted by OpenAI, has led to a growing tension between corporate profitability and the social responsibility to ensure that AI serves the broader public good. This tension is particularly problematic in high-stakes situations, such as civil defense, where access to advanced AI technologies could mean the difference between life and death. Balancing these competing interests is essential to ensure that AI remains aligned with humanity's best interests and does not exacerbate existing inequalities by leaving vulnerable populations without access to critical technology.

Balancing Corporate Profitability with Social Responsibility: The Role of AI in the Public Sphere

The rise of AI as a powerful tool for addressing complex problems—from climate change to global pandemics—has placed it squarely within the public sphere. The implications of AI stretch far beyond the profit margins of private companies, affecting public health, safety, and social welfare on a global scale. As such, the ethical responsibility of corporations that develop and deploy AI technologies must extend beyond mere profitability. In the case of OpenAI, which has transitioned from a non-profit organization to a capped-profit model, the need to balance corporate interests with social responsibility is particularly pronounced.

a. Corporate Profitability vs. Social Responsibility

- The development of AI is expensive, and companies like OpenAI have legitimate reasons to seek funding and generate revenue to sustain their research and development efforts. However, when the pursuit of profit becomes the primary driver of decision-making, it can lead to significant ethical dilemmas. AI models like GPT-4, which have the potential to transform civil defense and disaster response, are being restricted by paywalls and subscription fees, making them inaccessible to many of the communities that need them the most.
- In the public sphere, AI should be viewed not merely as a commercial product but as a tool for addressing societal challenges. Just as governments regulate access to essential goods and services (such as healthcare, education, and utilities) to ensure they are available to all, there is a growing argument that AI—especially in life-saving applications—should be treated as a public good. This would require corporations like

OpenAI to embrace social responsibility as a core component of their business models, ensuring that their technologies are accessible to all, not just those who can afford them.

b. Public Good vs. Private Control

- Historically, technologies with the potential to benefit the public—such as vaccines, clean energy innovations, and the internet—have required a balance between private innovation and public access. While corporations have played a key role in developing these technologies, their widespread adoption has often relied on public funding, government intervention, or open-source models that prioritize accessibility over profit.
- In the case of AI, the stakes are particularly high. The role of AI in civil defense, for example, is not a niche application; it is a matter of public safety. AI tools that can predict natural disasters, coordinate emergency responses, and offer medical advice in times of crisis are too critical to be left behind closed doors. Corporations that develop such technologies must recognize their broader responsibility to society and prioritize making these tools available to those who need them most, even if doing so requires a shift away from purely profit-driven models.

Ensuring that AI Development Remains Aligned with Humanity's Best Interests, Especially in High-Stakes Situations Like Civil Defense

One of the most pressing ethical challenges in the development of AI is ensuring that it remains aligned with humanity's best interests. AI has the potential to offer transformative benefits in areas like healthcare, education, and civil defense, but if its development is driven solely by market forces, there is a risk that it will be deployed in ways that prioritize corporate interests over the public good. This is particularly concerning in high-stakes situations, such as civil defense, where the decisions made by AI developers can have life-or-death consequences.

a. High-Stakes Applications of AI

- In civil defense, the applications of AI are critical for public safety. AI tools can be used to predict natural disasters like hurricanes and earthquakes, optimize emergency response efforts, manage resources during crises, and provide medical guidance in areas where healthcare infrastructure is

overwhelmed or nonexistent. These high-stakes applications require a level of ethical scrutiny that goes beyond typical market considerations.

- AI developers must ensure that their technologies are not only effective but also accessible, transparent, and accountable. In situations where lives are at risk, it is not enough to simply develop cutting-edge technology; developers must also consider how that technology will be deployed, who will have access to it, and whether it will exacerbate existing inequalities. AI models like GPT-4, if restricted by high costs and proprietary controls, could lead to situations where wealthy communities have access to life-saving technologies while poorer or more vulnerable populations are left behind.

b. Ethical AI Development Frameworks

- There are several emerging frameworks for ethical AI development that aim to address these concerns. For example, the "AI for Good" initiative, led by organizations like the United Nations, advocates for the development of AI technologies that prioritize social welfare and address global challenges. Similarly, the Partnership on AI, a consortium of AI developers, researchers, and civil society organizations, has developed guidelines for ensuring that AI is developed in ways that are ethical, transparent, and aligned with public interest.
- These frameworks emphasize the importance of transparency, accountability, and inclusivity in AI development. Transparency ensures that AI models are open to scrutiny and can be understood by users, developers, and policymakers. Accountability requires that developers take responsibility for the social and ethical implications of their technologies, including ensuring that they are not deployed in ways that harm vulnerable populations. Inclusivity involves ensuring that AI tools are accessible to all, regardless of socioeconomic status or geographic location, particularly in areas like civil defense where access to technology can save lives.

How Greed-Driven Models Could Exacerbate Inequity and Leave Vulnerable Populations Without Access to Critical Technology

The profit-driven approach to AI development, as exemplified by OpenAI’s shift toward a capped-profit model, has the potential to exacerbate existing

inequalities and leave vulnerable populations without access to critical technologies. When AI models are controlled by corporations that prioritize profitability over accessibility, the benefits of these technologies are likely to be concentrated among wealthier individuals and organizations, while poorer or more marginalized communities are left behind.

a. The Digital Divide

- The digital divide—the gap between those who have access to modern information and communication technologies and those who do not—has long been a concern for policymakers, technologists, and civil society. The development of AI introduces a new layer of complexity to this issue. As AI becomes an increasingly important tool for everything from healthcare to education to civil defense, access to AI technologies is becoming a key determinant of social and economic opportunity.
- Greed-driven models that restrict access to AI tools through high pricing and proprietary controls risk deepening the digital divide. Wealthier communities and organizations that can afford to pay for AI services will benefit from the advanced capabilities of models like GPT-4, while poorer communities, particularly in developing countries or under-resourced regions, will be left without access to these technologies. This inequality in access could lead to further disparities in public health, education, and disaster resilience, as vulnerable populations are unable to deploy the AI tools that could help them prepare for and respond to crises.

b. Exacerbating Global Inequality

- The global inequality in access to AI tools is particularly concerning in the context of civil defense. Countries and regions that are most vulnerable to natural disasters, conflicts, and pandemics are often those with the least access to advanced technologies. For example, small island nations that are at risk of being devastated by climate change may lack the resources to deploy AI models that could help predict and mitigate the effects of rising sea levels or extreme weather events. Similarly, conflict zones where healthcare infrastructure has been destroyed may not have access to AI tools that could provide remote medical guidance or optimize resource distribution during humanitarian crises.

- By restricting access to AI models like GPT-4 through cloud-based services and subscription fees, companies like OpenAI are effectively locking out the communities that could benefit most from these technologies. This creates a vicious cycle in which vulnerable populations are left increasingly at risk, while wealthier regions and organizations continue to reap the benefits of AI-driven innovation.

c. The Ethical Implications of Exclusion

- The exclusion of vulnerable populations from access to AI tools is not only an ethical issue but also a strategic one. In a globalized world, the failure to ensure that AI is accessible to all can have far-reaching consequences. Natural disasters, pandemics, and conflicts do not respect national borders, and the inability of one region to respond effectively to a crisis can have ripple effects across the globe. Ensuring that AI is developed and deployed in ways that are inclusive and equitable is not just a matter of justice; it is also a matter of global security and resilience.
- Greed-driven models that prioritize corporate profitability over social responsibility risk undermining the very purpose of AI: to solve complex problems and improve the human condition. By leaving vulnerable populations without access to critical technologies, these models fail to realize the full potential of AI as a force for good. To address this, AI developers must embrace ethical frameworks that prioritize inclusivity and ensure that the benefits of AI are shared broadly, rather than concentrated among the wealthy and powerful.

Conclusion

As AI becomes increasingly integral to civil defense, disaster response, and public safety, it is essential that its development and deployment are guided by ethical and strategic considerations that prioritize the public good. The tension between corporate profitability and social responsibility must be carefully managed to ensure that AI remains aligned with humanity's best interests, particularly in high-stakes situations like civil defense. Greed-driven models that restrict access to AI through high pricing and proprietary controls risk exacerbating existing inequalities and leaving vulnerable populations without the tools they need to prepare for and respond to crises. To prevent this, AI developers must embrace

ethical frameworks that prioritize transparency, accountability, and inclusivity, ensuring that AI serves the needs of all people, not just those who can afford it.

Conclusion

As we move deeper into the age of artificial intelligence, the ethical and strategic decisions we make about how these technologies are developed, distributed, and applied will have profound implications for the future of global resilience, equality, and public safety. AI models like GPT-4 represent some of the most advanced tools we have for addressing the challenges of the modern world, from natural disasters to public health crises to civil defense. Yet, the current profit-driven model adopted by companies like OpenAI is preventing many communities, particularly those most vulnerable to crises, from accessing these life-saving technologies. This is not merely a technological limitation but an ethical failure to ensure that AI serves the common good.

Reaffirming the Need for AI Models Like GPT-4 to Be Freely Available for the Public Good

The fundamental promise of artificial intelligence is its potential to improve the human condition. In high-stakes situations, such as disaster preparedness and civil defense, AI can save lives by providing timely, data-driven insights, optimizing resource allocation, and offering real-time medical guidance. However, this potential is being undermined by the current business models that restrict access to these powerful technologies through high costs and proprietary control. For AI to fulfill its promise as a public good, models like GPT-4 must be made freely available to the communities and individuals who need them most.

The argument for open access is particularly compelling in crisis scenarios. In times of natural disaster, conflict, or public health emergencies, the ability to deploy AI locally—without reliance on external cloud services or costly subscription models—could dramatically improve the speed and effectiveness of the response. Whether it's using AI to predict the path of a hurricane, manage the distribution of emergency supplies, or guide non-experts through life-saving medical procedures, access to AI tools can be the difference between chaos and coordination, between life and death. By making AI models freely available, we

can ensure that every community has the tools it needs to prepare for, respond to, and recover from crises.

Moreover, the need for local, decentralized AI deployment is more urgent than ever as we face an increasing frequency of global crises. Climate change is accelerating the occurrence of extreme weather events, geopolitical instability is heightening the risk of conflict, and public health threats continue to emerge in unpredictable ways. In this context, AI should be treated not as a luxury for the few but as a critical tool for the many. Open-source models that can be downloaded, customized, and deployed locally would empower communities to become more self-reliant and resilient in the face of these challenges.

A Call for OpenAI and Other AI Developers to Prioritize Human Welfare Over Profit

To make this vision a reality, companies like OpenAI must return to their founding principles. When OpenAI was established, it did so with the mission of ensuring that artificial intelligence would be developed for the benefit of all humanity, with an emphasis on openness, transparency, and accessibility. However, the organization's shift toward a profit-driven model has betrayed this mission, placing financial interests above the public good. OpenAI and other AI developers must recognize the immense responsibility they hold and prioritize human welfare over profit.

This means embracing a model of open access and distribution that aligns with the original vision of AI as a democratizing force for good. OpenAI has the power to lead the way by making its most advanced models, like GPT-4, freely available to communities around the world, particularly those that are most vulnerable to crises. By doing so, the organization could reclaim its place as a leader in ethical AI development, setting an example for the industry and ensuring that AI is used not to deepen inequalities but to bridge them.

Moreover, this call for open access extends beyond just OpenAI. The entire AI development ecosystem must adopt a more socially responsible approach to technology creation and dissemination. Governments, non-profits, academic institutions, and private companies all have roles to play in ensuring that AI is not monopolized by a few wealthy corporations but is instead shared as a global resource. This shift would require a fundamental rethinking of how AI is funded,

developed, and distributed, but it is a shift that is necessary if we are to build a more just and resilient world.

Final Thoughts on the Long-Term Benefits of Open-Source AI in Building a Safer, More Resilient World

The long-term benefits of open-source AI are vast and transformative. By making AI models freely available and allowing communities to customize them for their unique needs, we unlock a wave of innovation that extends far beyond what any single corporation or government could achieve. Open-source AI fosters collaboration, transparency, and creativity, as communities, researchers, and developers work together to solve the world's most pressing challenges. It removes the barriers of cost and access, ensuring that everyone, regardless of location or economic status, has the tools to build a safer and more resilient future.

Open-source AI also strengthens local capacity. When communities can deploy AI models on local hardware, they reduce their reliance on centralized systems that may be vulnerable to disruption during crises. This decentralized approach allows communities to take ownership of their own disaster preparedness and response efforts, tailoring AI tools to their specific circumstances and needs. Over time, this leads to stronger, more self-sufficient communities that are better equipped to handle the unpredictable challenges of the future.

Additionally, open-source AI has the potential to close the gap between the Global North and Global South, between wealthier regions and those that are often left behind in the march of technological progress. By making AI tools available to all, we create a more equitable distribution of knowledge and resources, empowering communities that have historically been excluded from the benefits of technological innovation. This is not just a matter of fairness; it is a matter of global security and stability. In a world where crises are increasingly global in scope, ensuring that every community has the tools to respond effectively is in the best interest of all.

In conclusion, the time has come to rethink how we develop and distribute AI. The current profit-driven model is unsustainable, inequitable, and ultimately harmful to the broader goal of building a safer, more resilient world. By embracing open-source AI, we can unlock the full potential of this transformative

technology, ensuring that it serves humanity as a whole rather than a privileged few. It is not too late for OpenAI and other developers to return to their original mission and lead the charge toward a future where AI is freely available, empowering communities around the world to face the challenges of the 21st century with confidence and resilience.

Epilogue: GPT-4o in the Aftermath

In the distant aftermath of a nuclear war, the world as we once knew it has changed irrevocably. The internet, once a ubiquitous force connecting people and resources globally, is most certainly down. Communication networks have collapsed, power grids have failed, and centralized systems are no longer functioning. Amid this chaos, an isolated community—cut off from the outside world—relies on a lifeline of technology that was once considered cutting-edge but is now the foundation of their survival.

Running on a $2,000 Seeed-NVIDIA Jetson Orin, GPT-4o—a locally deployed, optimized version of GPT-4—has become an essential tool for this community. The Orin device, a compact yet powerful piece of hardware, sits quietly in a small bunker, humming away as it processes requests, provides guidance, and answers questions with unparalleled precision. The world outside may have been reduced to rubble, but within this community, GPT-4o has become a beacon of knowledge in an otherwise desolate landscape.

The absence of the internet has not diminished GPT-4o's ability to function. In fact, its strength lies in its local deployment. Without the need for external servers or cloud-based processing, the model is fully operational, running directly on the community's hardware. It offers everything from real-time medical advice to survival strategies, translating its vast database of pre-loaded knowledge into practical, life-saving applications.

Medical Advice

The makeshift clinic in the center of the community is often crowded with people seeking care. With doctors in short supply, GPT-4o has taken on the role of a virtual medical advisor. Residents describe their symptoms to the AI, and GPT-4o provides step-by-step instructions on how to treat common injuries, illnesses, and

conditions. Whether it's administering first aid for burns, diagnosing infections, or advising on the proper use of scavenged medical supplies, GPT-4o's knowledge is a critical asset. It has even been programmed to offer psychological support, helping residents cope with the trauma of their new reality.

Resource Management

Managing resources in a post-apocalyptic world is one of the most difficult tasks facing the community. Food, water, fuel, and medicine are all in short supply, and ensuring that these limited resources are distributed fairly and efficiently requires careful planning. GPT-4o assists local leaders by analyzing data on stockpiles and usage patterns, providing recommendations on how best to ration supplies and extend the community's survival window. The AI can project future needs, estimate the impact of various resource allocation strategies, and offer insights into sustainable practices that can be implemented with the available materials.

Agricultural Guidance

With access to industrial food production and supply chains lost, the community has turned to growing its own food. But farming in the irradiated soil is no easy task. GPT-4o has become an agricultural advisor, helping the community identify the most viable crops to grow in harsh conditions, guiding them on the best planting techniques, and advising on soil remediation using the few remaining tools and chemicals available. The AI provides a detailed understanding of crop cycles, water conservation methods, and pest control, helping the community establish a fragile but critical food supply chain.

Disaster Mitigation

Though the worst has already happened, the world remains volatile. Weather patterns have become erratic, and radiation levels fluctuate unpredictably. GPT-4o helps monitor local radiation levels using simple, locally built Geiger counters, predicting when it's safest to venture outside for supplies or farming. It also uses stored weather data to advise the community on how to protect themselves from the increasingly unpredictable elements, whether it be extreme cold, heat, or radiation storms.

Education and Communication

Despite the bleakness, the community's leadership understands that knowledge is their most valuable resource. GPT-4o becomes a teacher, educating the next generation in essential survival skills, mathematics, basic engineering, and even literature. Although the internet is gone, GPT-4o's database includes a vast array of texts that can be accessed by the community's youth, ensuring that education continues despite the collapse of modern civilization. Furthermore, GPT-4o helps facilitate communication within the community, offering translation services for the diverse languages spoken by the community members who came together from different regions before the war.

A New Kind of Resilience

In this isolated, post-nuclear world, GPT-4o represents more than just an advanced piece of technology—it embodies the resilience of the human spirit. The decision to deploy AI locally, without reliance on external systems, was a strategic one made before the war, and it is now paying dividends. The community has embraced this decentralized model, where knowledge and survival do not depend on fragile global infrastructures but instead on the powerful synergy between human ingenuity and artificial intelligence.

The irony is not lost on the community's elders, who recall a time when technology like GPT-4 was controlled by powerful corporations, locked behind paywalls, and accessible only to the wealthy. Now, in the most unexpected of circumstances, this technology has become a shared resource, accessible to all who remain. In the face of unimaginable destruction, the shift toward open-source and local deployment has given this small group of survivors a fighting chance.

In the end, GPT-4o has become a testament to what is possible when technology is made freely available for the public good. Its presence in the bunker is a reminder that, while the internet may be gone, knowledge persists, and with it, hope. The community continues to endure, its survival made possible by the foresight to deploy AI locally—an investment in resilience that is now proving its worth, day after day, in a world where survival itself is a daily struggle.

And so, in this new world, GPT-4o is not just a tool; it is a lifeline, a silent partner in humanity's fight to rebuild, one answer at a time.

References

1. **OpenAI**. (2023). *The evolution of OpenAI's mission: Balancing public benefit and commercial interests*. Retrieved from https://openai.com
2. **EleutherAI**. (2021). *GPT-J and GPT-Neo: Open-source large language models*. EleutherAI Organization. Retrieved from https://www.eleuther.ai
3. **BigScience Project**. (2022). *BLOOM: A multilingual open-source language model*. Retrieved from https://bigscience.huggingface.co
4. **Partnership on AI**. (2021). *Ethical frameworks for AI development*. Partnership on AI. Retrieved from https://www.partnershiponai.org
5. **United Nations AI for Good Initiative**. (2022). *AI for Good: Harnessing the power of artificial intelligence for public benefit*. United Nations. Retrieved from https://www.itu.int/en/ITU-T/AI/Pages/default.aspx
6. **Hugging Face**. (2023). *Hugging Face: Open-source models and AI tools for developers*. Hugging Face. Retrieved from https://huggingface.co
7. **Jetson Orin**. (2023). *NVIDIA Jetson Orin: Compact edge AI devices for powerful computing*. NVIDIA Corporation. Retrieved from https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin
8. **Algorithmic Justice League**. (2021). *Ensuring fairness and transparency in AI development*. Retrieved from https://www.ajl.org
9. **OpenMRS**. (2020). *Open-source medical record system for developing countries*. Retrieved from https://openmrs.org
10. **OpenSSL**. (2023). *Open-source cryptography toolkit for securing digital communications*. Retrieved from https://www.openssl.org
11. **Wireshark**. (2022). *Open-source network protocol analyzer*. Retrieved from https://www.wireshark.org
12. **World Health Organization**. (2022). *Using AI for global health security: The role of open-source technologies*. Retrieved from https://www.who.int

LOCALLY DEPLOYED
ORIN
GPT JETSON 4o
GPT-4o